ABSTRACT

A technique for managing a customer's traffic in a network node, such as a service provider edge device, involves dedicating a group of queues in the network node to the customer, performing queue-specific rate shaping on the customer's traffic according to queue-specific bandwidth limitations, and performing group-specific rate shaping on the customer's traffic as a whole according to a group-specific bandwidth limitation. In an embodiment, the queues in the group are associated with different types of customer traffic in order to provide type-specific rate shaping. Further, the queues may be prioritized among each other such that unused excess bandwidth is distributed among the different traffic types in priority order.